LEARNING FEATURES OF MUSIC FROM SCRATCH

John Thickstun[1], **Zaid Harchaoui**[2] **& Sham Kakade**[1,2]
[1] Department of Computer Science and Engineering, [2] Department of Statistics
University of Washington
Seattle, WA 98195, USA
{thickstn,sham}@cs.washington.edu, zaid@uw.edu

arXiv:1611.09827v1 [stat.ML] 29 Nov 2016

ABSTRACT

We introduce a new large-scale music dataset, MusicNet, to serve as a source of supervision and evaluation of machine learning methods for music research. MusicNet consists of hundreds of freely-licensed classical music recordings by 10 composers, written for 11 instruments, together with instrument/note annotations resulting in over 1 million temporal labels on 34 hours of chamber music performances under various studio and microphone conditions.

We define a multi-label classification task to predict notes in musical recordings, along with an evaluation protocol. We benchmark several machine learning architectures for this task: i) learning from "hand-crafted" spectrogram features; ii) end-to-end learning with a neural net; iii) end-to-end learning with a convolutional neural net. We show that several end-to-end learning proposals outperform approaches based on learning from hand-crafted audio features.

1 INTRODUCTION

Music research has benefited recently from the effectiveness of machine learning methods on a wide range of problems from music recommendation (van den Oord et al., 2013; McFee & Lanckriet, 2011) to music generation (Driedger et al., 2015); see also the recent demos of the Google Magenta project[1]. As of today, there is no large publicly available labeled dataset for the simple yet challenging task of note prediction for classical music. The MIREX MultiF0 Development Set (Benetos & Dixon, 2011) and the Bach10 dataset (Duan et al., 2011) together contain less than 7 minutes of labeled music. These datasets were designed for method evaluation, not for training supervised learning methods.

This situation stands in contrast to other application domains of machine learning. For instance, in computer vision, large labeled datasets such ImageNet (Russakovsky et al., 2015) were fruitfully used to train end-to-end learning architectures. Learned feature representations have outperformed traditional hand-crafted low-level visual features and lead to tremendous progress for image classification. In (Humphrey et al., 2012), Humphrey, Bello, and LeCun issued a call to action: "Deep architectures often require a large amount of labeled data for supervised training, a luxury music informatics has never really enjoyed. Given the proven success of supervised methods, MIR would likely benefit a good deal from a concentrated effort in the curation of sharable data in a sustainable manner."

We introduce here a new large labeled dataset, MusicNet, that we make publicly available[2] to foster progress learning feature representations of music. MusicNet is a large corpus of aligned labels on freely-licensed classical music recordings, made possible by licensing initiatives of the European Archive, the Isabella Stewart Gardner Museum, Musopen, and various individual artists. The dataset consists of 34 hours of human-verified aligned recordings, containing a total of $1,299,329$ individual labels on segments of these recordings. Table 1 summarizes statistics of MusicNet.

This work was supported by the Washington Research Foundation Fund for Innovation in Data-Intensive Discovery, and the program "Learning in Machines and Brains" (CIFAR).

[1] https://magenta.tensorflow.org/

[2] http://homes.cs.washington.edu/~thickstn/musicnet.html.

MusicNet

Minutes	Labels	Recordings	Error Rate
2,048	1,299,329	330	4.0%

Ensemble	Minutes	Labels
Solo Piano	917	576,471
String Quartet	405	259,702
Accompanied Violin	148	124,886
Piano Quartet	73	60,362
Accompanied Cello	63	37,557
String Sextet	48	33,248
Piano Trio	46	28,873
Piano Quintet	25	27,545
Wind Quintet	43	24,820
Horn Piano Trio	30	18,799
Wind Octet	23	14,635
Clarinet-Cello-Piano Trio	25	13,447
Pairs Clarinet-Horn-Bassoon	24	12,218
Clarinet Quintet	26	11,184
Solo Cello	49	10,876
Accompanied Clarinet	20	10,049
Solo Violin	30	8,837
Violin and Harpsichord	16	7,469
Viola Quintet	15	4,156
Solo Flute	8	2,214

Composer	Minutes	Labels
Beethoven	1,085	736,072
Schubert	253	146,648
Brahms	192	133,109
Mozart	156	99,641
Bach	184	62,782
Dvorak	56	46,261
Cambini	43	24,820
Faure	33	22,349
Ravel	27	21,243
Haydn	15	6,404

Instrument	Minutes	Labels
Piano	1346	794,532
Violin	874	230,484
Viola	621	99,407
Cello	800	99,132
Clarinet	173	24,426
Bassoon	102	14,954
Horn	132	11,468
Oboe	66	8,696
Flute	69	8,310
Harpsichord	16	4,914
String Bass	38	3,006

	Piano	Violin	Cello	Viola	Clarinet	Bassoon	Horn	Oboe	Flute	Bass	Harpsichord
Notes	83	51	51	51	41	36	41	28	37	43	51

Table 1: Summary statistics of the MusicNet dataset.

The focus of this paper is the problem of learning low-level features of music from raw audio data. We define a multi-label classification task to predict notes in musical recordings, along with an evaluation protocol. We benchmark a variety of machine learning architectures for this task: i) learning from "hand-crafted" spectrogram features; ii) end-to-end learning with a neural net; iii) end-to-end learning with a convolutional neural net. We show that several end-to-end learning architectures outperform approaches based on learning from hand-crafted audio features. The experimental results suggest that, for each of the proposed models, modulated sine-like waveform features are stable, optimal low-level features of musical audio. The learned low-level features are visualized in Figure 1.

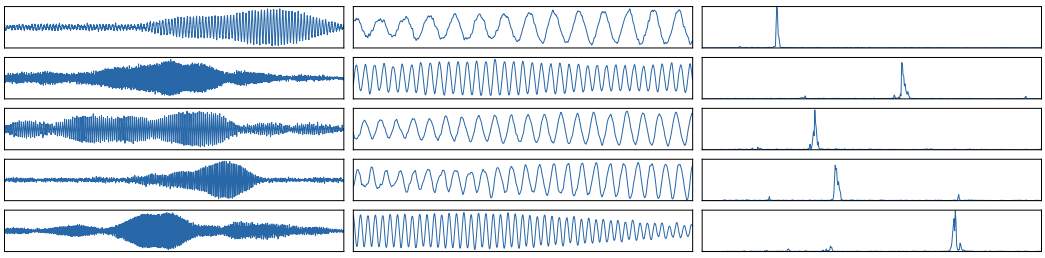

Figure 1: (Left) Bottom-level weights learned by a two-layer ReLU network trained with ℓ_2 regularized ($\lambda = 1$) square loss for multi-label classification on raw audio recordings. (Middle) Magnified view of the center of each set of weights. (Right) The spectrogram of each set of weights.

2 MUSICNET

MusicNet is a large collection of freely-licensed recordings together with labels on these recordings exemplified in Table 2. We find that large amounts of data are essential to recovering useful features from music; see Sect. 4.1 for details. The Lakh dataset, released this summer based on the work of Raffel & Ellis (2015), offers note-level annotations for many 30-second clips of pop music in the Million Song Dataset (McFee et al., 2012). Other large-scale music databases are less useful for supervised representation learning. The RWC dataset (Goto et al., 2003) does not have note-level labels. The MAPS dataset (Emiya et al., 2010) consists of synthesized data, which expressive models could overfit. The Mazurka project[3] consists of commercial music; accessing this dataset comes at a cost and inconvenience, requiring researchers to track down a multitude of commercial recordings. Both the MAPS and Mazurka datasets are comprised entirely of piano music.

The MusicNet dataset consists of 330 recordings of a variety of instruments arranged in small chamber ensembles under various studio and microphone conditions. The recordings average 6 minutes in length. The shortest recording in the dataset is 55 seconds and the longest is almost 18 minutes. Table 1 summarizes the statistics of MusicNet with breakdowns into various types of labels. Table 2 demonstrates examples of labels from the MusicNet dataset.

Start	End	Instrument	Note	Measure	Beat	Note Value
45.29	45.49	Violin	G5	21	3	Eighth
48.99	50.13	Cello	A#3	24	2	Dotted Half
82.91	83.12	Viola	C5	51	2.5	Eighth

Table 2: MusicNet labels on the Pascal String Quartet's recording of Beethoven's Opus 127, String Quartet No. 12 in E-flat major, I - Maestoso - Allegro. Creative commons use of this recording is made possible by the work of the European Archive.

MusicNet labels come from 513 label classes using the most naive definition of a class: distinct instrument/note combinations. The breakdowns reported in Table 1 indicate the number of distinct notes that appear for each instrument in our dataset. For example, while a piano has 88 keys only 83 of them are performed in MusicNet. For many tasks a note's value will be a part of its label, in which case the number of classes will expand by approximately an order of magnitude after taking the cartesian product of the set of classes with the set of values: quarter-note, eighth-note, triplet, etc. We also remark that labels regularly overlap in the time series creating polyphonic multi-labels.

MusicNet is heavily skewed towards Beethoven, thanks to the composer's popularity among performing ensembles. The dataset is also skewed towards Solo Piano due to an abundance of digital scores available for piano works. For training purposes, we expect that researchers may want to augment this dataset to increase coverage of instruments such as Flute and Oboe that are underrepresented in MusicNet. Researchers who do not need to distribute their dataset can make use of immense libraries of commercial recordings. These recordings can be labeled using the alignment protocol described in Sect. 3.

3 DATASET CONSTRUCTION

We have collected 158 hours of freely-licensed classical music recordings from the European Archive, the Isabella Stewart Gardner Museum, Musopen, and various artists' collections. We have also collected 1,618 digital scores in the MIDI format from online resources including the Classical Archives (`classicalarchives.com`) Suzuchan's Classic MIDI (`suzumidi.com`) and HarfeSoft (`harfesoft.de`). We can produce an alignment in cases where a digital score in our collection corresponds to a freely-licensed recording. In addition to our aligned scores, we have gathered MIDI scores containing an additional $6,550,760$ labels; we make these labels available to researchers who wish to augment MusicNet with commercial recordings.

Music-to-score alignment is a long-standing problem in the music research and signal processing communities (Raphael, 1999). Dynamic time warping (DTW) is a classical approach to this prob-

[3] `http://www.mazurka.org.uk/`

lem. An early reference using DTW is Orio & Schwarz (2001) where music is aligned to a crude synthesis of the score designed to capture some of the structure of an overtone series. We make use of side information from a synthesizer, aligning music to an artificial performance of a score. To the best of our knowledge, commercial synthesis was first used for the purpose of alignment in Turetsky & Ellis (2003).

The majority of previous work on alignment focuses on pop music. This is more challenging than aligning classical music because commercial synthesizers do a poor job reproducing the wide variety of vocal and instrumental timbers that appear in modern pop. Furthermore, pop features anharmonic instruments such as drums for which natural metrics on frequency representations–including ℓ^2–are unmeaningful. We find that a variant of the techniques described in Turetsky & Ellis (2003) works robustly for classical music to score alignment; we discuss our evaluation of this procedure and its error rate on MusicNet in the appendix.

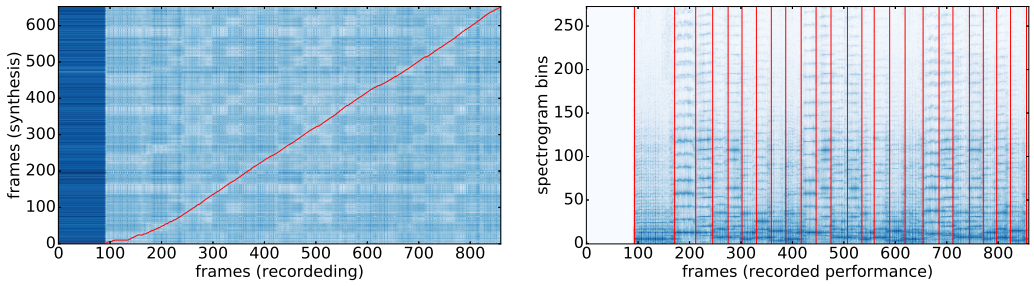

Figure 2: (Left) Heatmap visualization of local alignment costs between the synthesized and recorded spectrograms, with the optimal alignment path in red. The block from $x = 0$ to $x = 100$ corresponds to silence at the beginning of the recorded performance. The slope of the alignment can be interpreted as an instantaneous tempo ratio between the recorded and synthesized performances. The curvature in the alignment between $x = 100$ and $x = 175$ corresponds to an extension of the first notes by the performer. (Right) Annotation of note onsets on the spectrogram of the recorded performance, determined by the alignment shown on the left.

In order to align the performance with a score, we need to define a metric that compares short segments of the score with segments of a performance. Musical scores can be expressed as binary vectors in $E \times K$ where $E = \{1, \ldots, n\}$ and K is a dictionary of notes. Performances reside in $\mathbb{R}^{T \times p}$, where $T \in \{1, \ldots, m\}$ is a sequence of time steps and p is the dimensionality of the spectrogram at time T. Given some local cost function $C : (\mathbb{R}^p, K) \to \mathbb{R}$, a score $\mathbf{Y} \in E \times K$, and a performance $\mathbf{X} \in \mathbb{R}^{T \times p}$, the alignment problem is to

$$\begin{aligned} \underset{t \in \mathbb{Z}^n}{\text{minimize}} \quad & \sum_{i=1}^{n} C(\mathbf{X}_{t_i}, \mathbf{Y}_i) \\ \text{subject to} \quad & t_0 = 0, \\ & t_n = m \\ & t_i \leq t_j \qquad \text{if } i < j. \end{aligned} \tag{1}$$

Dynamic time warping gives an exact solution to the problem in $\mathcal{O}(mn)$ time and space.

The success of dynamic time warping depends on the metric used to compare the score and the performance. Previous works can be broadly categorized into three groups that define an alignment cost C between segments of music $\mathbf{x}$ and score $\mathbf{y}$ by injecting them into a common normed space via maps Ψ and Φ:

$$C(\mathbf{x}, \mathbf{y}) = \|\Psi(\mathbf{x}) - \Phi(\mathbf{y})\| \tag{2}$$

The most popular approach–which we have adopted–maps the score into the space of the performance (Orio & Schwarz, 2001; Turetsky & Ellis, 2003; Soulez et al., 2003). An alternative approach maps both the score and performance into some third space, commonly a chromogram space (Hu et al., 2003; Izmirli & Dannenberg, 2010; Joder et al., 2013). Finally, some recent methods consider alignment in score space, taking $\Phi = \mathrm{Id}$ and learning Ψ (Garreau et al., 2014; Lajugie et al., 2016).

With reference to the general cost (2), we must specify the maps Ψ, Φ, and the norm $\|\cdot\|$. We compute the cost in the performance feature space $\mathbb{R}^p$, hence we take $\Psi = \mathrm{Id}$. For our features, we use the log-spectrogram with a window size of 2048 samples. We use a stride of 512 samples between features. Hence adjacent feature frames are computed with 75% overlap. For audio sampled at 44.1kHz, this results in a feature representation with $44,100/512 \approx 86$ frames per second. A discussion of these parameter choices can be found in the appendix. The map Φ is computed by a synthetizer: we used Plogue's Sforzando sampler together with Garritan's Personal Orchestra 4 sample library.

For a (pseudo)-metric on $\mathbb{R}^p$, we take the ℓ^2 norm $\|\cdot\|_2$ on the low 50 dimensions of $\mathbb{R}^p$. Recall that $\mathbb{R}^p$ represents Fourier components, so we can roughly interpret the k'th coordinate of $\mathbb{R}^p$ as the energy associated with the frequency $k \times (22,050/1024) \approx k \times 22.5$Hz, where $22,050$Hz is the Nyquist frequency of a signal sampled at 44.1kHz. The 50 dimension cutoff is chosen empirically: we observe that our alignments are much more accurate using a small number of low-frequency bins rather than the full space $\mathbb{R}^p$. Synthesizers do not accurately reproduce the high-frequency features of a musical instrument; by ignoring the high frequencies, we align on a part of the spectrum where the synthesis is most accurate. Our choice of cutoff is aggressive compared to usual settings; for instance, Turetsky & Ellis (2003) propose cutoffs in the 2.5kHz range. The fundamental frequencies of many notes in our dataset are higher than the $50 \times 22.5\text{Hz} \approx 1\text{kHz}$ cutoff. Nevertheless, we find that all notes align well using only the low-frequency information.

4 METHODS

We consider identification of notes in a segment of audio $\mathbf{x} \in \mathcal{X}$ as a multi-label classification problem, modeled as follows. Assign each audio segment a binary label vector $\mathbf{y} \in \{0,1\}^{128}$. The 128 dimensions correspond to frequency codes for notes, and $\mathbf{y}_n = 1$ if note n is present at the midpoint of $\mathbf{x}$. Let $f : \mathcal{X} \to \mathcal{H}$ indicate a feature map. We train a multivariate linear regression to predict $\hat{\mathbf{y}}$ given $f(\mathbf{x})$, which we optimize for square loss. The vector $\hat{\mathbf{y}}$ can be interpreted as a multi-label estimate of notes in $\mathbf{x}$ by choosing a threshold c and predicting label n iff $\hat{\mathbf{y}}_n > c$. We search for c on a sampled subset of MusicNet, optimizing for F-score with grid search.

4.1 RELATED WORK

Learning on raw audio has been considered in both the music and speech communities. Supervised learning on music has been driven by access to labeled datasets. Pop music annotations with chord labels (Harte, 2010) have lead to a long line of work on supervised chord recognition, most recently Korzeniowsk & Widmer (2016). Song-level genre labels and various other metadata have also attracted substantial work on representation learning; a recent example is Choi et al. (2016). There is also substantial work modeling raw audio representations of speech; a current example is Tokuda & Zen (2016).

Because access to large labeled datasets was historically limited, much of the work in the music community is unsupervised. Variants of non-negative matrix factorization are popular in the music information retrieval community, for example Khlif & Sethu (2015). Berg-Kirkpatrick et al. (2014) develops a Bayesian model for piano music. Recent work from Google DeepMind explores generative models of raw audio, including music (van den Oord et al., 2016).

4.2 MULTI-LAYER PERCEPTRONS

We construct a two-layer ReLU network using the features $f_i(\mathbf{x}) = \max(0, w_i^T \mathbf{x})$. Figure 1 illustrates a selection of weights w_i learned by the bottom layer of this network, optimized for multi-label classification using square loss. The weights learned by the network are modulated sinusoids. This explains the effectiveness of spectrograms and related transforms as a low-level representation of musical audio. The weights decay at the boundaries, analogous to Gabor filters in vision. This behavior is explained by our labeling methodology: the audio segments used here are approximately $1/3$ of a second long, and a segment is given a note label if that note is on in the center of the segment. Therefore information at the boundaries of the segment is less useful for prediction than information nearer to the center.

4.3 SPECTROGRAMS

Spectrograms are an engineered feature representation for musical audio signals, available in popular software packages such as librosa (McFee et al., 2015). Spectrograms are closely related to the two-layer ReLU network discussed above. If $\mathbf{x} = (x_1, \ldots, x_t)$ denotes a segment of an audio signal of length t then we can define

$$\text{Spec}_k(\mathbf{x}) \equiv \left| \sum_{s=1}^{t} e^{iks} x_s \right|^2 = \left(\sum_{s=1}^{t} \cos(ks) x_s \right)^2 + \left(\sum_{s=1}^{t} \sin(ks) x_s \right)^2 .$$

These features are not precisely learnable by the two-layer ReLU network. But recall that $|x| = \max(0, x) + \max(0, -x)$ and if we take weight vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^T$ with $u_t = \cos(kt)$ and $v_t = \sin(kt)$ then the ReLU network can learn

$$f_{k,\cos}(\mathbf{x}) + f_{k,\sin}(\mathbf{x}) = |\mathbf{u}^T \mathbf{x}| + |\mathbf{v}^T \mathbf{x}| = \left| \sum_{s=1}^{t} \cos(ks) x_s \right| + \left| \sum_{s=1}^{t} \sin(ks) x_s \right| .$$

We call this family of features a ReLUgram and observe that it has a similar form to the spectrogram; we merely replace the $x \mapsto x^2$ non-linearity of the spectrogram with $x \mapsto |x|$. These features achieve similar performance to spectrograms on our classification task (see Table 3).

4.4 WINDOW SIZE

When we parameterize a network, we must choose the width of the set of weights in the bottom layer. This width is called the receptive field in the vision community; in the music community it is called the window size. Traditional frequency analyses, including spectrograms, are highly sensitive to the window size. Windows must be long enough to capture relevant information, but not so long that they lose temporal resolution; this is the classical time-frequency tradeoff. Furthermore, windowed frequency analysis is subject to boundary effects, known as spectral leakage. Classical signal processing attempts to dampen these effects with hand-crafted window functions, which apply a mask that attenuates the signal at the boundaries (Rabiner & Schafer, 2007).

Our models learn good window functions. If we parameterize our models with a large window size then the model will learn that distant information is irrelevant to local prediction, so the magnitude of the learned weights will attenuate at the boundaries (see Figure 1). We therefore focus our attention on two window sizes: 2048 samples, which captures the local content of the signal, and 16,384 samples, which is sufficient to capture almost all relevant context (again we refer to Figure 1; substantially larger window sizes would be a needless computation burden, because the weights at further distances will approximately vanish).

4.5 REGULARIZATION

The size of MusicNet is essential to achieving the results in Figure 1. Prior work on end-to-end audio learning was unable to recover clean sinusoidal features from data (Dieleman & Schrauwen, 2014). We encountered similar problems when optimizing on a small subset of MusicNet. In Figure 3 (Left) we optimize a two-layer ReLU network on $65,000$ monophonic data points; compare this to similar results in Figure 3 of Dieleman & Schrauwen (2014). We can recover sinusoidal features on the small dataset using heavy regularization, but this destroys classification performance; regularizing with dropout poses a similar tradeoff. By contrast, Figure 3 (Right) shows weights learned on the full MusicNet dataset using no regularization whatsoever. We are still exploring the effects of ℓ_2 regularization on the full dataset; preliminary experiments suggest that a modest amount of regularizer stabilizes the optimization and produces even cleaner features without sacrificing performance.

4.6 CONVOLUTIONAL NETWORKS

Previously, we estimated $\hat{\mathbf{y}}$ by regressing against $f(\mathbf{x})$. We now consider a convolutional model that regresses against features of a collection of shifted segments $\mathbf{x}_\ell$ near to the original segment $\mathbf{x}$. The learned features of this network are visually comparable to those learned by the fully connected network (Figure 1). We have experimented with the stride and number of convolutions in this network.

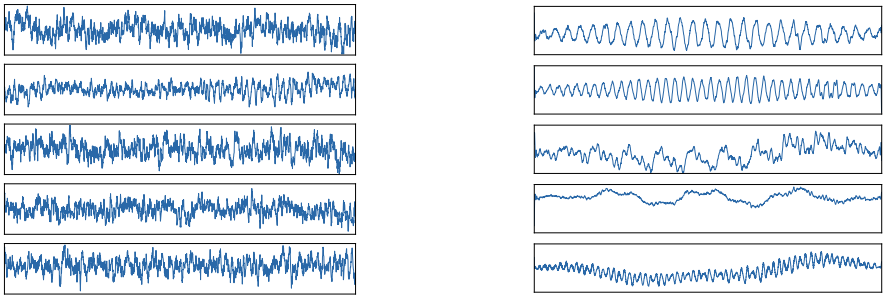

Figure 3: (Left) Features learned by a 2-layer ReLU network trained on small monophonic subset of MusicNet. (Right) Features learned by the same network, trained on the full MusicNet dataset.

The results reported in Table 3 were achieved using a 64-sample stride and 97 convolutions across a window of $16,384$ samples, using a receptive field of $10,240$ samples. Performance correlates with the resolution of the stride and the number of convolutions, but the learned features are consistent across parameterizations. We also experimented with average and max pooling operations. In all cases the learned features are comparable to those of a fully connected network.

5 RESULTS

We hold out a test set of 3 recordings for all the results reported in this section:

- Bach's Prelude in D major for Solo Piano. WTK Book 1, No 5. Performed by Kimiko Ishizaka. MusicNet recording id 2303.
- Mozart's Serenade in E-flat major. K375, Movement 4 - Menuetto. Performed by the Soni Ventorum Wind Quintet. MusicNet recording id 1819.
- Beethoven's String Quartet No. 13 in B-flat major. Opus 130, Movement 2 - Presto. Released by the European Archive. MusicNet recording id 2382.

Our test set is a representative sampling of MusicNet: it covers most of the instruments in the dataset in small, medium, and large ensembles. The test data points are evenly spaced segments separated by 512 samples, between the 1st and 91st seconds of each recording. For the wider features, there is substantial overlap between adjacent segments. Each segment is labeled with the notes that are on in the middle of the segment.

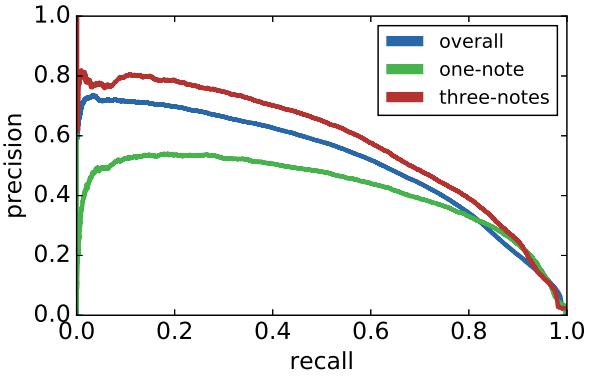

Figure 4: Precision-recall curves for the convolutional network on the test set. Curves are evaluated on subsets of the test set consisting of all data points (blue); points with exactly one label (monophonic; green); and points with exactly three labels (red).

We evaluate our models on three scores: precision, recall, and average precision. The precision score is the count of correct predictions by the model (across all data points) divided by the total number

of predictions by the model. The recall score is the count of correct predictions by the model divided by the total number of (ground truth) labels in the test set. Precision and recall are parameterized by the note prediction threshold c (see Sect. 4). By varying c, we construct precision-recall curves (see Figure 4). The average precision score is the area under the precision-recall curve.

Model	Features	Precision	Recall	Average Precision
Linear	512-point spectrogram	22.1%	47.0%	22.0%
Linear	1024-point spectrogram	28.9%	52.5%	30.2%
Linear	1024-point ReLUgram	24.1%	60.0%	29.3%
Linear	4096-point spectrogram	35.2%	63.2%	40.3%
Linear	8192-point spectrogram	32.1%	65.6%	37.5%
MLP, 500 nodes	2048 raw samples	37.5%	57.5%	41.0%
MLP, 2500 nodes	2048 raw samples	40.5%	58.6%	43.8%
AvgPool, 5 stride	2048 raw samples	38.9%	59.0%	43.3%
MLP, 500 nodes	16384 raw samples	40.0%	63.7%	45.5%
CNN, 64 stride	16384 raw samples	43.2%	70.7%	52.0%

Table 3: Benchmark results on MusicNet for models discussed in this paper. All models were optimized using the Tensorflow library (Abadi et al.). The MLP is a 2-layer ReLU network with an unregularized square loss objective. The AvgPool model is parameterized by 500 hidden nodes and 11 convolutions. The CNN was parameterized with 500 hidden nodes and 97 convolutions. We report the precision and recall corresponding to the best F_1-score.

A spectrogram of length n is computed from $2n$ samples, so the linear 1024-point spectrogram model is directly comparable to the MLP runs with 2048 raw samples. We find that our learned features[4] significantly beat the performance of spectrograms. Our discussion of windowing in Sect. 4.4 partially explains this. Figure 5 suggests a second reason. Recall (Sect. 4.3) that the spectrogram features can be interpreted as the magnitude of the signal's inner product with sine waves of linearly spaced frequencies. In contrast, our networks learn weights with frequencies distributed similarly to the distribution of notes in our dataset (Figure 5). This gives our network higher resolution in the most critical frequency regions.

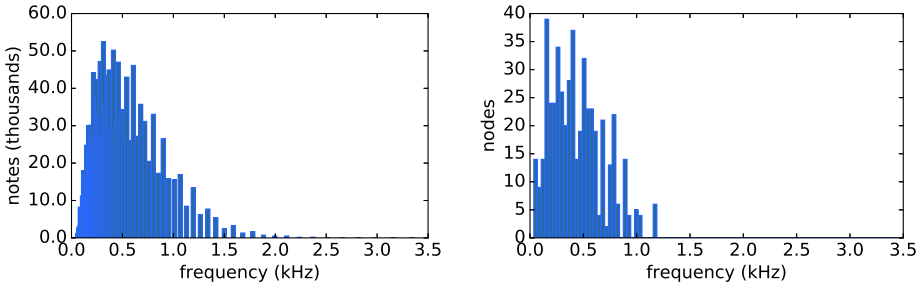

Figure 5: (Left) The frequency distribution of notes in MusicNet. (Right) The frequency distribution of learned nodes in a 500-node, two-layer ReLU network.

In future work, we plan to investigate learned mid-level and high-level features of musical audio. While mid-level features could capture harmonic structure, high-level features could capture the overall structure of a recording. Both mid-level and high-level representations require the low-level features learned in this paper as building blocks to extract short-term and long-term memory temporal structures.

[4]A demonstration using learned MLP features to synthesize a musical performance is available on the author's webpage at `http://homes.cs.washington.edu/~thickstn/demos.html`

REFERENCES

M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems. URL `http://tensorflow.org/`.

E. Benetos and S. Dixon. Joint multi-pitch detection using harmonic envelope estimation for polyphonic music transcription. *IEEE Selected Topics in Signal Processing*, 2011.

T. Berg-Kirkpatrick, J. Andreas, and D. Klein. Unsupervised transcription of piano music. *NIPS*, 2014.

K. Choi, G. Fazes, and M. Sandler. Automatic tagging using deep convolutional neural networks. *ISMIR*, 2016.

S. Dieleman and B. Schrauwen. End-to-end learning for music audio. *ICASSP*, 2014.

J. Driedger, T. Prätzlich, and M. Müller. Let It Bee - Towards NMF-inspired audio mosaicing. *ISMIR*, 2015.

Z. Duan, B. Pardo, and C. Zhang. Multiple fundamental frequency estimation by modeling spectral peaks and non-peak regions. *TASLP*, 2011.

V. Emiya, R. Badeau, and B. David. Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle. *TASLP*, 2010.

D. Garreau, R. Lajugie, S. Arlot, and F. Bach. Metric learning for temporal sequence alignment. *NIPS*, 2014.

M. Goto, H. Hashiguchi, T. Nishimura, and R. Oka. RWC music database: Music genre database and musical instrument sound database. *ISMIR*, 2003.

C. Harte. *Towards Automatic Extraction of Harmony Information from Music Signals*. PhD thesis, Department of Electrical Engineering, Queen Mary, University of London, 2010.

N. Hu, R. B. Dannenberg, and G. Tzanetakis. Polyphonic audio matching and alignment for music retrieval. *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 2003.

E. J. Humphrey, J. P. Bello, and Y. LeCun. Moving beyond feature design: Deep architectures and automatic feature learning in music informatics. *ISMIR*, 2012.

O. Izmirli and R. B. Dannenberg. Understanding features and distance functions for music sequence alignment. *ISMIR*, 2010.

C. Joder, S. Essid, and G. Richard. Learning optimal features for polyphonic audio-to-score alignment. *TASLP*, 2013.

A. Khlif and V. Sethu. An iterative multi range non-negative matrix factorization algorithm for polyphonic music transcription. *ISMIR*, 2015.

F. Korzeniowsk and G. Widmer. Feature learning for chord recognition: the deep chroma extractor. *ISMIR*, 2016.

R. Lajugie, P. Bojanowski, P. Cuvillier, S. Arlot, and F. Bach. A weakly-supervised discriminative model for audio-to-score alignment. *ICASSP*, 2016.

B. McFee and G. Lanckriet. Learning multi-modal similarity. *JMLR*, 2011.

B. McFee, T. Bertin-Mahieux, D. P. W. Ellis, and G. Lanckriet. The million song dataset challenge. *Proceedings of the 21st International Conference on World Wide Web*, 2012.

B. McFee, C. Raffel, D. Liang, D. P. W. Ellis, M. McVicar, E. Battenberg, and O. Nieto. librosa: Audio and music signal analysis in python. *SCIPY*, 2015.

N. Orio and D. Schwarz. Alignment of monophonic and polyphonic music to a score. *International Computer Music Conference*, 2001.

L. Rabiner and R. Schafer. Introduction to digital speech processing. *Foundations and trends in signal processing*, 2007.

C. Raffel and D. P. W. Ellis. Large-scale content-based matching of MIDI and audio files. *ISMIR*, 2015.

C. Raphael. Automatic segmentation of acoustic musical signals using hidden markov models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1999.

O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *IJCV*, 2015.

F. Soulez, X. Rodet, and D. Schwarz. Improving polyphonic and poly-instrumental music to score alignment. *ISMIR*, 2003.

K. Tokuda and H. Zen. Directly modeling voiced and unvoiced components in speech waveforms by neural networks. *ICASSP*, 2016.

R. J. Turetsky and D. P. W. Ellis. Ground-truth transcriptions of real music from force-aligned midi syntheses. *ISMIR*, 2003.

A. van den Oord, S. Dieleman, and B. Schrauwen. Deep content-based music recommendation. *NIPS*, 2013.

A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu. WaveNet: A generative model for raw audio. *arXiv preprint*, 2016.

A VALIDATING THE MUSICNET LABELS

We validate the aligned MusicNet labels with a listening test. We create an aural representation of an aligned score-performance pair by mixing a short sine wave into the performance with the frequency indicated by the score at the time indicated by the alignment. We can listen to this mix and, if the alignment is correct, the sine tones will exactly overlay the original performance; if the alignment is incorrect, the mix will sound dissonant.

We have listened to a substantial portion of each recording in the aligned dataset: the beginning, several random samples of middle, and the end. Any mixes with substantially incorrect alignments were rejected from the dataset. Failed alignments were mostly attributable to mismatches between the midi and the recording. The most common reason for rejection was musical repeats. Classical music often contains sections with indications that they be repeated a second time; in classical music performance culture, it is often considered acceptable to ignore these directions. If the score and performance make different choices regarding repeats, a mismatch arises. When the score omits a repeat that occurs in the performance, the alignment will typically warp over the entire repeated section, with correct alignments before and after. When the score includes an extra repeat, the alignment typically compresses it into very short segment, with correct alignments on either side. We rejected alignments exhibiting either of these issues from the dataset.

From the aligned performances that we deemed sufficiently accurate to admit to the dataset, we also randomly sampled 30 clips for more careful annotation and analysis. We weighted the sample to cover a wide coverage of recordings with various instruments, ensemble sizes, and durations. For each sampled performance, we randomly selected a 30 second clip. Using software transforms, it is possible to slow a recording down to approximately 1/4 speed. Two of the clips were too richly structured and fast to precisely analyze (slowing the signal down any further introduces artifacts that make the signal difficult to interpret). Even in these two rejected samples, the alignments sound substantially correct.

For the other 28 clips, we carefully analyzed the aligned performance mix and annotated every alignment error. Two of the authors are classically trained musicians: we independently checked for errors and we our analyses were nearly identical. Where there was disagreement, we used the more pessimistic author's analysis. Note that we do not catch every type of error: we are likely to miss performance mistakes that maintain the meter of the performance, but for professional recordings such mistakes are rare. Over our entire set of clips we averaged a 4.0% error rate.

We can also qualitatively characterize the types of errors we observed. The most common types of errors are anticipations and delays: a single, or small sequence of labels is aligned to a slightly early or late location in the time series. Another common source of error is missing ornaments and trills: these are short flourishes in a performance are sometimes not annotated in our score data, which results in a missing annotation in the alignment. Finally, there are rare performance errors in the recordings and transcription errors in the score.

B ALIGNMENT PARAMETER ROBUSTNESS

Our definitions of audio featurization and the alignment cost function were contingent on several parameter choices. These choices were optimized by systematic exploration of the parameter space: we investigated what happens as we vary each parameter and made the choices that gave the best results in our listening tests. The bottom line is that there is no magic in our parameter choices: choosing the parameters carefully yields marginal gains, but alignment performance degrades gracefully as the choices diverge from the optimum.

The quality of alignments improve uniformly with the quality of synthesis. The time-resolution of labels improves uniformly as the stride parameter decreases; minimization of stride is limited by system memory constraints. We find that the precise phase-invariant feature specification has very little effect on alignment quality. We experimented with spectrograms and log-spectrograms using windowed and un-windowed signals. Alignment quality is largely unaffected.

The other parameters are governed by a tradeoff curve; the optimal choice is determined by balancing desirable outcomes. The fourier window size is a classic tradeoff between time and frequency resolution. The ℓ^2 norm can be understood as a tradeoff between the extremes of ℓ^1 and ℓ^∞. The ℓ^1

norm is too egalitarian: the preponderance of errors due to synthesis quality add up and overwhelm the signal. On the other hand, the ℓ^∞ norm ignores too much of the signal in the spectrogram. The spectrogram cutoff, discussed extensively in Sect. 3, is also a tradeoff between synthesis quality and maximal use of information

C PRECISION & RECALL CURVES

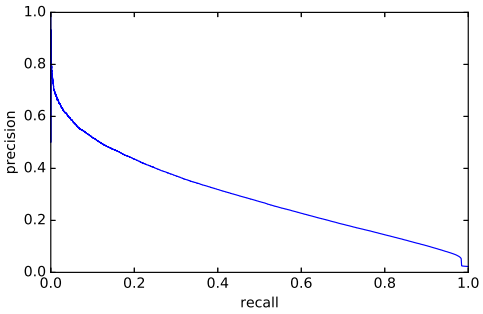

Figure 6: The linear ReLUgram model.

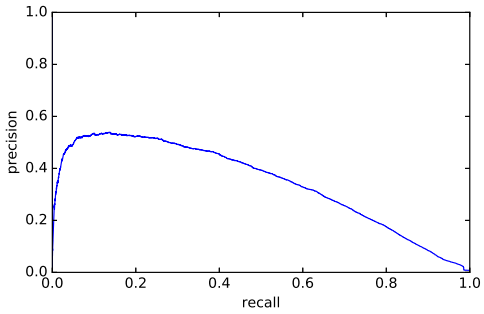

Figure 7: The 500 node, 2048 raw sample MLP.

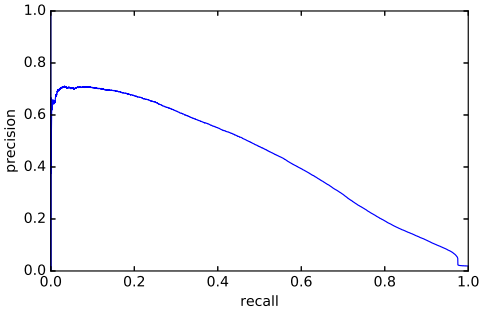

Figure 8: The 2500 node, 2048 raw sample MLP.

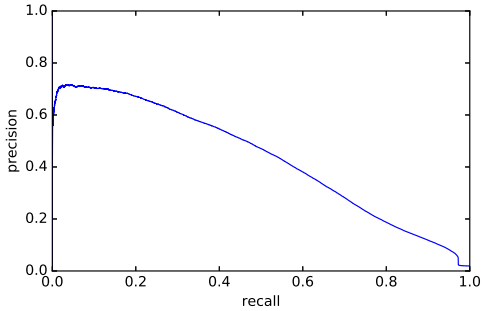

Figure 9: The average pooling model.

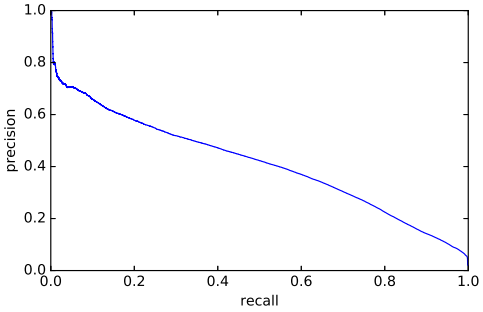

Figure 10: The 500 node, 16384 raw sample MLP ($\lambda = 1$).

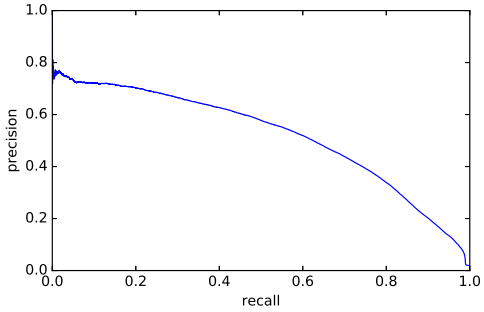

Figure 11: The convolutional model.

D ADDITIONAL RESULTS

We report additional results on splits of the test set described in Sect. 5.

Model	Features	Precision	Recall	Average Precision
MLP, 500 nodes	2048 raw samples	64.1%	45.6%	53.3%
MLP, 2500 nodes	2048 raw samples	53.5%	68.7%	63.8%
AvgPool, 5 stride	2048 raw samples	50.9%	70.2%	63.1%
CNN, 64 stride	16384 raw samples	56.3%	77.5%	70.3%

Table 4: The Soni Ventorum recording of Mozart's Wind Quintet K375 (MusicNet id 1819).

Model	Features	Precision	Recall	Average Precision
MLP, 500 nodes	2048 raw samples	31.1%	34.2%	23.9%
MLP, 2500 nodes	2048 raw samples	23.8%	53.9%	26.8%
AvgPool, 5 stride	2048 raw samples	23.2%	53.3%	26.1%
CNN, 64 stride	16384 raw samples	29.4%	69.3%	37.4%

Table 5: The European Archive recording of Beethoven's String Quartet No. 13 (MusicNet id 2382).

Model	Features	Precision	Recall	Average Precision
MLP, 500 nodes	2048 raw samples	70.7%	31.7%	47.4%
MLP, 2500 nodes	2048 raw samples	54.0%	50.4%	50.8%
AvgPool, 5 stride	2048 raw samples	50.4%	50.3%	49.6%
CNN, 64 stride	16384 raw samples	51.0%	62.2%	54.4%

Table 6: The Kimiko Ishizaka recording of Bach's Prelude in D major (MusicNet id 2303).